

# Market-Based Multirobot Coordination: A Survey and Analysis

M. Bernardine Dias, Robert Zlot, Nidhi Kalra, and Anthony Stentz

CMU-RI-TR-05-13

April 2005

Robotics Institute  
Carnegie Mellon University  
Pittsburgh, Pennsylvania 15213

© Carnegie Mellon University

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>APR 2005</b>		2. REPORT TYPE		3. DATES COVERED <b>00-00-2005 to 00-00-2005</b>	
4. TITLE AND SUBTITLE <b>Market-Based Multirobot Coordination: A Survey and Analysis</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Carnegie Mellon University ,Robotics Institute,Pittsburgh,PA,15213</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release; distribution unlimited</b>					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT <b>Market-based multirobot coordination approaches have received significant attention and gained considerable popularity within the robotics research community in recent years. They have been successfully implemented in a variety of domains ranging from mapping and exploration to robot soccer. The research literature on market-based approaches to coordination has now reached a critical mass that warrants a survey and analysis. This paper addresses this need by providing an introduction to market-based multirobot coordination, a comprehensive review of the state of the art in the field, and a discussion of remaining challenges.</b>					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT  <b>Same as Report (SAR)</b>	18. NUMBER OF PAGES  <b>25</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			



## Abstract

Market-based multirobot coordination approaches have received significant attention and gained considerable popularity within the robotics research community in recent years. They have been successfully implemented in a variety of domains ranging from mapping and exploration to robot soccer. The research literature on market-based approaches to coordination has now reached a critical mass that warrants a survey and analysis. This paper addresses this need by providing an introduction to market-based multirobot coordination, a comprehensive review of the state of the art in the field, and a discussion of remaining challenges.

---

<sup>1</sup>A version of this paper will appear in the Proceedings of the IEEE – Special Issue on Multi-Robot Systems (expected publication in 2006).



# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Overview</b>	<b>2</b>
2.1	Definition of a Market-based Approach . . . . .	2
2.2	The Range of Coordination Approaches . . . . .	2
<b>3</b>	<b>Planning</b>	<b>4</b>
3.1	Related Work . . . . .	4
3.1.1	Planning and Task Allocation . . . . .	4
3.1.2	Planning and Task Decomposition . . . . .	5
3.1.3	Planning and Task Execution . . . . .	5
3.2	Future Challenges . . . . .	5
<b>4</b>	<b>Dynamic Events and Environments</b>	<b>5</b>
4.1	Related Work . . . . .	6
4.1.1	Robustness and Fluidity . . . . .	6
4.1.2	Response Speed . . . . .	6
4.1.3	Online Tasks . . . . .	7
4.1.4	Uncertainty . . . . .	7
4.2	Future Challenges . . . . .	7
<b>5</b>	<b>Quality of Solution</b>	<b>7</b>
5.1	Related Work . . . . .	7
5.1.1	Practical Approaches . . . . .	7
5.1.2	Theoretical Guarantees . . . . .	8
5.2	Future Challenges . . . . .	8
<b>6</b>	<b>Scalability</b>	<b>9</b>
6.1	Related Work . . . . .	9
6.1.1	Computation and Communication Considerations . . . . .	9
6.1.2	Opportunistic Centralization . . . . .	9
6.2	Future Challenges . . . . .	10
<b>7</b>	<b>Heterogeneous Teams</b>	<b>10</b>
7.1	Related Work . . . . .	10
7.2	Future Challenges . . . . .	11
<b>8</b>	<b>Tight Coordination</b>	<b>11</b>
8.1	Related Work . . . . .	11
8.2	Future Challenges . . . . .	12
<b>9</b>	<b>Learning and Adaptation</b>	<b>12</b>
9.1	Related Work . . . . .	12
9.2	Future Challenges . . . . .	12

<b>10 Generality</b>	<b>13</b>
10.1 Related Work . . . . .	13
10.1.1 Flexibility . . . . .	13
10.1.2 Extensibility . . . . .	13
10.1.3 Implementation . . . . .	13
10.1.4 Comparisons . . . . .	14
10.2 Future Challenges . . . . .	14
<b>11 Conclusions and Future Directions</b>	<b>14</b>

# 1 Introduction

As robots become an integral part of human life they are charged with increasingly difficult tasks. Many of these tasks can be better achieved by a team of robots than by a single robot. By working together, robots can complete tasks faster, increase system robustness, improve solution quality, and achieve tasks impossible for a single robot. Nevertheless, coordinating such a team requires overcoming many formidable research challenges.

Given a team of robots, a limited amount of resources, and a team task, researchers must develop a method of distributing the resources among the team so the task is accomplished well, even as teammates' interactions, the environment, and the mission change. Humans have dealt with similar problems for thousands of years with increasingly sophisticated market economies in which the individual pursuit of profit leads to the redistribution of resources and an efficient production of output. The principles of a market economy can be applied to multirobot coordination. In this virtual economy, the robots are traders, tasks and resources are traded commodities, and virtual money acts as currency. Robots compete to win tasks and resources by participating in auctions that produce efficient distributions based on specified preferences. When the system is appropriately designed, each robot acts to maximise its individual profit and simultaneously improves the efficiency of the team.

This paper is motivated by the growing popularity of market-based multirobot coordination approaches and the lack of a comprehensive review of these approaches. As contributors to pioneering this research area, the authors have several years of experience in designing and implementing market-based coordination approaches for multi-robot systems. In this paper the authors draw upon their collective research experience to review the current literature and recommend future research directions. A more detailed analysis of the relevant literature is available as a technical report [24]. This paper makes three significant contributions to the robotics literature. First, it introduces market based approaches by discussing the motivating philosophy, defining the requirements of such approaches, analyzing their strengths and weaknesses, and placing them appropriately in the context of the larger set of approaches to multirobot coordination. Second, this paper surveys and analyses the relevant literature. Finally, it inspires and directs future research on this topic through a discussion of remaining challenges.

The scope of this paper is limited to market-based approaches for coordinating teams that include robots. Furthermore, this review principally considers approaches that actively reason about the existence of other agents when coordinating the team, in contrast to approaches in which agents coexist. Nevertheless, related work on market-based multiagent research in software agent domains and other relevant publications are included as necessary to augment discussion.

The following section provides an introduction to market-based mechanisms for readers less familiar with the field and a quick reference for researchers designing and implementing market-based coordination mechanisms. This overview is followed by a extensive review of market-based multirobot coordination approaches categorized and analysed across several relevant dimensions: planning, dynamic events and environments, solution quality, scalability, heterogeneity, tight coordination, learning and adaptation, and generality. The paper concludes with a summary of the survey and



future challenges in this research area.

## 2 Overview

The earliest examples of market-based multi-agent coordination appeared in the literature over twenty-five years ago [14, 40] and have been modified and adopted for multirobot coordination in more recent years. These approaches have been successfully implemented in a variety of domains.

### 2.1 Definition of a Market-based Approach

Most “market-based” multirobot and multiagent coordination approaches share a set of underlying elements. Market theory provides precise definitions for several of these elements. Borrowing from both bodies of literature, we define a market-based multirobot coordination approach based on the following requirements:

- The team is given an objective that can be decomposed into subcomponents that are achievable by individuals or subteams. To solve the problem, the team has at its disposal a limited set of resources that is distributed among the team members.
- A global objective function quantifies the system designer’s preferences for all possible solutions.
- An individual utility function specified for each robot quantifies that robot’s preferences for its individual resource usage and contributions towards the team objective. Evaluating this function cannot require global or perfect information about the state of the team or team objective. Subteam preferences can also be quantified through a combination of individual utilities.
- A mapping is defined between the team objective function and individual and subteam utilities. This mapping addresses how the individual production and consumption of resources and individuals’ advancement of the team objective affect the overall solution.
- Resources and individual or subteam objectives can be redistributed using a mechanism such as an auction. This mechanism accepts as input teammates’ utilities and computes an outcome that maximizes the utility of the agent controlling the mechanism. In a well-designed mechanism, maximizing the controlling agent’s utility results in improving the team objective function value.

Thus, a multirobot coordination approach that satisfies all of the above requirements can generally be considered as a market-based approach.

### 2.2 The Range of Coordination Approaches

In virtually all robotic application domains, generating optimal solutions in a computationally tractable manner is highly advantageous. Unfortunately, these objectives

conflict in multirobot systems where optimal coordination is typically  $\mathcal{NP}$ -hard . The challenges are compounded by team requirements that include operation in dynamic environments, inconsistent information, unreliable and limited communication, interaction with humans, and system failures. A spectrum of coordination approaches has emerged to negotiate these demands.

At one end of the spectrum, fully centralized approaches employ a single agent to coordinate the entire team. In theory, this agent can produce optimal solutions by gathering all relevant information and planning for the entire team. In reality, fully centralized approaches are rarely tractable for large teams, can suffer from a single point of failure, have high communication demands, and are usually sluggish to respond to changes perceived by distributed sources. Thus, centralized approaches are best suited for applications where teams are small and the environment is static or global state information is easily available. At the other end of the spectrum, in fully-distributed systems, robots rely solely on local knowledge. Such approaches are typically very fast, flexible to change, and robust to failures but can produce highly suboptimal solutions since good local solutions may not sum to a good global solution. Applications where large teams carry out relatively simple tasks with no strict requirements for efficiency are best served by fully distributed coordination schemes. A vast majority of coordination approaches have elements that are centralized and distributed and thus reside in the middle of the spectrum. Market-based approaches fall into this hybrid category, and, if designed well, they can opportunistically adapt to dynamic conditions to produce more centralized or more distributed solutions.

Market-based approaches effectively meet the practical demands of robot teams while producing efficient solutions by capturing the respective strengths of both distributed and centralized approaches. First, they can distribute much of the planning and execution over the team and thereby retain the benefits of distributed approaches, including robustness, flexibility, and speed [44, 45, 18, 13]. They also have elements of centralized systems to produce better solutions. Auctions concisely gather information about the team and distribute resources in a team-aware context. Researchers have also developed methods of opportunistically centralizing markets [12, 23] and providing guarantees of solution quality [28]. However, market-based approaches are not without their weaknesses. In domains where fully centralized approaches are feasible, market-based approaches can be more complex to implement and produce poorer solutions. In domains where fully distributed approaches suffice, market-approaches can be unnecessarily complex in design and can require excessive communication and computation. In other domains, the communication demands of auctions may be infeasible. Perhaps the biggest drawbacks of market-based approaches currently are the lack of performance guarantees and the lack of formalization in designing appropriate cost and revenue functions to capture domain requirements.

The following sections discuss market-based multirobot coordination in greater detail along the dimensions mentioned in the introduction. Each section introduces the topic and its challenges, defines the goals and appropriate evaluation metrics, reviews the relevant literature, and identifies the remaining research challenges.

## 3 Planning

In multirobot teams, planning can be used to coordinate robots to accomplish the team mission. Unfortunately, optimal planning for multirobot systems is typically an  $\mathcal{NP}$ -hard problem. The challenge then is to have tractable planning that produces good solutions. Market-based approaches manage this by distributing planning over the entire team to produce solutions quickly. When required or when resources permit, markets can behave more centrally and plan over larger portions of the team to improve solution quality. We evaluate approaches by how well the local planning captures the complexity of the problem while still being tractable.

### 3.1 Related Work

#### 3.1.1 Planning and Task Allocation

Task allocation is the problem of feasibly assigning a set of tasks to a team of robots in a way that optimizes the global objective function. Market-based approaches often distribute planning required for task allocation through the auction process: each robot locally plans the achievement of the offered tasks, computes its costs, and encapsulates the costs in its bids. We believe *TraderBots* [8] is most flexible in this respect. In addition to using local planning, *TraderBots* allows opportunistically centralized planning: when resources permit, “leaders” plan lower-cost redistributions of a subset of tasks to a subgroup of robots and improve the team solution. Other approaches are listed in Section 5.1.1. However, opportunistic centralized planning in market-based approaches is a relatively new topic that is yet to be formalized.

Related problems such as the allocation of constrained subtasks and the allocation of roles have somewhat different planning requirements:

**Allocating Constrained Subtasks** In many domains, tasks are temporally constrained with respect to each other. They may be partially ordered or may need to start or finish within a common time frame. During assignment, robots can incorporate the cost of meeting constraints into their bids [30]. More realistically, they must coordinate during execution to reschedule and accommodate team and task changes since the initial allocation [29, 19].

**Allocating Roles** In team games one usually assigns positions such as “primary offense” or “supporting defense” instead of tasks such as “shoot the ball” or “capture a rebound.” These positions can be classified as roles. More generally a role defines a collection of related actions or behaviors. Indeed, in many domains it is more natural to think of teammates playing roles than completing distinct tasks. In market-based approaches, role allocation uses the same auction-bid-award protocol as task allocation. However, robots generate bids by evaluating a fitness function that reflects how well the current state matches the requirements of the role. Once allocated, a robot locally plans the execution of actions and behaviors specified by its role. Market-based role allocation has primarily been demonstrated in the robot soccer domain (e.g. [26]).

### 3.1.2 Planning and Task Decomposition

Although many approaches to task allocation assume that a list of primitive or simple tasks is input to the system, a mission is often more naturally described at a more abstract level. In these cases, multirobot systems must also plan the decomposition of a mission into subtasks.

There are two common approaches to this planning problem. In the *decompose-then-allocate* method, a single agent recursively decomposes the task into simple subtasks which are allocated to the team [6]. In the *allocate-then-decompose* method, complex tasks are first allocated to robots; then, each robot locally decomposes its awarded tasks [5]. It is also possible to include instances of both techniques [19].

By decoupling the decomposition and allocation, these approaches do not consider the complete solution space and may find highly suboptimal solutions. Zlot and Stentz [44] address this issue by generalizing tasks to *task trees* within a peer-to-peer trading market. Robots may bid at any level of task abstraction in a tree with the option of redecomposing an awarded complex task more efficiently. Experiments in complex task domains have demonstrated that using task trees improves solution quality over two-stage approaches.

### 3.1.3 Planning and Task Execution

Significantly less work has been done to plan coordination between teammates during task execution. Nevertheless, it becomes necessary when robots interfere with each other during execution. Azarm and Schmidt [2] address collision avoidance of independent robots using market-based techniques. It is also necessary in domains where teammates continuously constrain each other's actions. Such coordination is more tightly-coupled and therefore difficult to distribute. In the Hoplites framework [23], constraints are captured in the utility function and planning is local when constraints are easily satisfied. When constraint satisfaction becomes difficult, teammates attempt resolution through negotiated participation in more centrally-developed joint plans. This is related to the idea of "leaders" in *TraderBots* [11].

## 3.2 Future Challenges

Replanning is a major remaining challenge and relates closely to the continuing challenges for teams operating in dynamic environments (Section 4.2). Existing market mechanisms are not yet fully capable of replanning task distributions, redecomposing tasks, rescheduling commitments, and replanning coordination during execution. Another important and relevant area of research is understanding the formation of subteams and enabling their positive interaction using market-based methods.

## 4 Dynamic Events and Environments

Operating in dynamic environments poses a variety of challenges. The principle challenges relevant to coordinating a team are ensuring graceful degradation of solution

quality with failures, enabling team functionality despite imperfect and uncertain information, maintaining high response speed to dynamic events, and effectively accommodating evolving task and team composition. Relevant evaluation metrics include the diversity of failures the team can accommodate, the required quantity and certainty in perceived information, the team's response speed to dynamic events, the fluidity of the team, and the overall solution quality produced by the team in the face of dynamic events.

## **4.1 Related Work**

### **4.1.1 Robustness and Fluidity**

Many multirobot applications require some level of robustness. Robotic systems are incredibly complex and their physical interactions with the environment make them highly prone to failures. Three principle categories of faults are communication failures, partial malfunctions, and robot death, all of which are gracefully handled by market-based approaches [13].

Teams can often perform more effectively if teammates can communicate [45]. However, communication failures occur in a variety of domains and range from occasional loss of messages to loss of all communication. If the team is informed of a common task, consequent disruptions in communication are gracefully handled by using opportunistic auctioning solely for improving solution quality [45, 13]. As with communication disruptions, partial malfunctions limit a robot's capability but retain the robot's planning facilities. Active reasoning about failed resources allows robots to reallocate tasks that they can no longer complete due to malfunctions [13], and monitoring progress in short-duration tasks allows detection and response to faults [17]. In the case of robot death, the affected robot cannot aid in the recovery process. However, robots can monitor heartbeats of teammates and re-auction the tasks previously assigned to the dead robot [13]. If malfunctioning robots can be repaired and return to the team, the coordination approach should accommodate both the exit and entrance of the repaired robots. Towing of disabled robots for repair [3] and re-entry of repaired robots [13] have also been demonstrated using market-based approaches. Additional fluidity can be demonstrated by changing the team composition during execution by adding a robot to the team [12].

### **4.1.2 Response Speed**

A key to successful task execution is often the ability to respond quickly to dynamic conditions. If information must always be channeled to another location for replanning, conditions can change too rapidly for the planner to keep up. Market-based approaches should be able to quickly respond to dynamic events since each robot can respond to local events independently without depending on a centralized coordinator to plan a response. However, the authors are unaware of any formalized study of response speed for any multirobot coordination approach.

#### **4.1.3 Online Tasks**

In many dynamic application domains, the demands on the robotic system can change during operation. Operators of a multirobot system may submit new tasks or alter or cancel existing tasks. Additionally, robots may generate new tasks during execution as they observe new information about their surroundings. Market-based approaches can often seamlessly incorporate online tasks by auctioning new tasks as they are introduced by an operator [8] or generated by the robots themselves [45, 21], or as they become available due to the completion of preceding tasks [5, 6, 18, 38]. Tasks can also be canceled during execution [12]. Some market-based approaches auction all tasks prior to execution and thus do not explicitly handle online tasks [4, 28, 31, 33].

#### **4.1.4 Uncertainty**

Most real-world multirobot applications require operation with only partial or changing information about the environment, the team, or the task. Fortunately, market-based approaches have few information requirements and can accommodate new information through frequent auctioning of tasks and resources. For instance, robots can begin tasks without prior map information [12] and dynamically reallocate tasks when new map information is gathered [42]. Robots also do not need to know the size or composition of the team to coordinate [18, 12]

### **4.2 Future Challenges**

While much can be done to improve market-based approaches to effectively operate in dynamic environments, a few key challenges are paramount. These challenges are incorporating contract breaches with appropriate penalties, developing more sophisticated methods for cooperative handling of partial malfunctions and repairs, and evaluating response speed and robustness to a variety of failures.

## **5 Quality of Solution**

One of the greatest strengths of a market is its ability to utilize the local information and preferences of its participants to arrive at an efficient solution given limited resources. The fundamental optimization problem encountered in market-based multirobot systems is the task allocation problem. Various global cost objectives appear in market-based systems [43], including the sum of individual robots' costs [4, 8, 20, 28, 31] and minimizing the maximum individual cost [33, 29].

### **5.1 Related Work**

#### **5.1.1 Practical Approaches**

Techniques used by market-based systems for task allocation can be broken down into two broad categories: one-task-per-robot (OTPR) approaches and sequencing approaches. In OTPR approaches, each robot behaves myopically and only considers

handling one task at any given time. In sequencing approaches, robots can order (or schedule [19, 29, 37]) a list of tasks and can therefore explicitly reason about the dependencies between multiple tasks and upcoming commitments.

The OTPR model arises in cases where each task requires an exclusive commitment of a robot (*e.g.* positional roles in robot soccer [26]). At the other extreme is continuous task allocation, by which we mean that the tasks being assigned are short-lived partial actions that bring the team goal closer to being realized. These actions can be updated as the environment changes or new observations are made, and then assigned to the robots by an auctioneer [18, 38]. OTPR approaches are also sometimes used for simplicity in cases where sequencing approaches could be used [6, 7, 17, 21]; however, this simplification is expected to result in less efficient solutions. In these systems, if there are more tasks than robots the remaining tasks can be assigned once robots become available.

Sequencing approaches have been used with centralized allocation, both for combinatorial auctions [4, 22, 31] and single task auctions [28, 33]. Combinatorial auctions have the potential to exploit the synergies between tasks, and thus produce better solutions as compared to single task auctions. Tovey *et al.* [43] show that for single-task auctions, choosing the appropriate bidding rule for a given global objective is important in terms of the resulting solution quality. Dias *et al.* [10] demonstrate that increasing the number of tasks awarded per auction can have a negative effect on the resulting solution quality, but requires less time to find a solution. Distributed or *peer-to-peer* allocation approaches allow auctions to take place between robots to improve an existing allocation [8, 10, 20, 29, 34]. Peer-to-peer trading can be viewed as a local search algorithm, and therefore is subject to local minima. By allowing exchanges of task sets of varying sizes [11], some local minima can be avoided [1, 35]. Additionally, in unknown or partially known environments where costs are constantly changing as new observations are made, peer-to-peer auctions act as a reallocation mechanism that can repair undesirable allocations.

### 5.1.2 Theoretical Guarantees

Lagoudakis *et al.* [27] provide a set of approximation bounds on solution quality for sequential single-task auctions for various team objective functions and bidding rules. For example, bidding marginal costs results in a *2-approximation* when minimizing a sum-of-costs team objective, while for a makespan objective the approximation ratio scales with the number of robots on the team. Sandholm proves that by using a sufficiently expressive set of contract types for peer-to-peer auctions the global minimum can be reached in a finite number of steps [35]. Gerkey presents some approximation results for OTPR models, including some online cases [17].

## 5.2 Future Challenges

While some theoretical guarantees for simple cases of auction algorithms are now known, implemented systems are generally more complex and can include online, multi-task, peer-to-peer, simultaneous, and overlapping auctions as well as task and

scheduling constraints. Additionally, solution quality depends on cost and utility estimates, which are sometimes difficult to accurately specify.

## 6 Scalability

A system is scalable if it can operate effectively even as the number of inputs or the size of inputs increases arbitrarily. The scalability of a multirobot coordination approach is typically evaluated by its ability to function as the team size or the task complexity increase. It is particularly desirable for approaches that are designed to be applicable to multiple domains, complex domains, and domains without upper bounds on team size. Scalability in market-based approaches may be limited by the computation and communication needs that arise from increasing auction frequency, bid complexity, and planning demands. However, market-based approaches can scale well in applications where the team mission can be decomposed into tasks that can be independently carried out by small sub-teams. An initial comparison of three approaches, including a market-based approach, on scalability is presented by Dias [8].

### 6.1 Related Work

#### 6.1.1 Computation and Communication Considerations

Single task auctions require a polynomial number of bid valuations and can be cleared in polynomial time. Additionally auction calls and bids require a polynomial number of messages to be sent. Combinatorial auctions, however, can in general require an exponential number of task bundles to be considered during bid valuation, and winner determination is  $\mathcal{NP}$ -hard. Additionally, an exponential number of messages would have to be sent to submit bids for the  $2^{|T|}$  possible task bundles. In order to make these problems tractable, the number of task bundles can be reduced either on the auctioneer side by offering only certain combinations of tasks [22], or on the bidder side by using heuristic task clustering algorithms [4, 11]. In some cases, due to the sparseness of the bids, the auction clearing step can be done optimally and quickly [36], although the resulting allocation is still likely to be suboptimal given that not all task bundles are considered.

Bid valuation itself may be computationally expensive. In many cases an  $\mathcal{NP}$ -hard problem must be solved in order to estimate task costs [4, 8, 20, 33, 34], or expensive task decompositions may be required at the bidding stage [44]. When there are many tasks to consider simultaneously – either from auctions offering many tasks [4, 8, 20, 29, 31, 44] or from many robots initiating auctions at roughly the same time [8, 44] – bidders can be overburdened with task valuation problems. As a result, they may not be able to meet auction deadlines or may tax their processors to the point of not being able to do any other useful work.

#### 6.1.2 Opportunistic Centralization

Opportunistic-centralization [8, 23] is a scalable way of incorporating the benefits of centralized planning into large market-based systems. In such approaches, centralized



resource or task allocation can be done over smaller subsets of tasks and team members as and when computational resources permit.

## 6.2 Future Challenges

Market-based approaches have yet to be implemented on teams of more than a few robots. Further work can be done to improve opportunistically centralized approaches to select subsets of tasks and team members more selectively to reduce unnecessary computation. Additionally, the problem of dealing with potential overflow in the amount of bid valuation computations is largely unaddressed.

# 7 Heterogeneous Teams

A heterogeneous team is composed of members who have different capabilities or play different roles. In contrast, homogeneous teams are composed of members with identical skills, or generalists, capable of all necessary tasks. Many complex tasks are achieved more effectively if decomposed into components that require different skills and executed by heterogeneous teams where members specialize at different tasks. Moreover, it is often simpler to design robots that specialize in a small set of skills than to design robots capable of all skills. Thus, heterogeneity is a highly desirable in many teams. However, the coordination problem is more difficult if the robots are heterogeneous. Nevertheless, a successful coordination approach will be able to accommodate any team regardless of its homogeneity or heterogeneity.

## 7.1 Related Work

A major challenge when dealing with heterogeneous teams is effective task or role allocation that optimizes the employment of individual skills. This requires reasoning about different robot models and capabilities. Market-based approaches are well-suited to coordinating heterogeneous teams because auctions can in some ways simplify the problem of reasoning about team skills. When a task or role is being auctioned, bids encapsulate individual ability to complete that task and can also encode the solution quality afforded by each member's resources, and even the opportunity cost of completing other tasks [37]. The auctioneer awards roles or tasks to team members according to the best bid without requiring knowledge of the team's capabilities. Thus, market-based approaches only require individual team members to recognize their own skills and resources and not those of teammates. An appropriately designed market can effectively coordinate heterogeneous teams. Coordinating a team with heterogeneous capabilities using a market-based approach has been demonstrated both on physical robots [39, 38, 17] and in simulation [37, 8]. Heterogeneous role allocation to a robot team using a market-based approach has also been demonstrated on both physical robots [17] and on simulated robot teams [16, 25]. However, appropriately designing the different components of a market-based approach to capture the complexities of a heterogeneous team executing a complex task can be difficult. One idea

proposed to deal with the problem of determining an appropriate currency in heterogeneous teams is to allow robots to swap tasks [20]. If robots can swap tasks, then every deal can be individually rational without needing to consider currency and payments. However, limiting contracts to swapping severely restricts the number of possible outcomes. Combining a market-based approach with the concept of team “plays” is another proposal for coordinating dynamic heterogeneous teams [9].

## 7.2 Future Challenges

Ultimately, coordination approaches must accommodate three levels of heterogeneity: heterogeneous robot teams, human-robot teams, and highly heterogeneous teams of humans, robots, and other agents. For market-based approaches, the challenges in this area include providing interfaces to allow better human participation, modeling of human preferences using appropriate reward functions, designing modular agents that can represent a diversity of team members, and implementing portable trading agents that can interface with a variety of planning and execution layers thus providing a plug-and-play capability for any team member to join the market.

# 8 Tight Coordination

Many tasks can be decomposed into subtasks that can be completed by individual robots. Consequently robots coordinate during task decomposition and allocation but not during execution. Such teams are *loosely coordinated*. On the other hand, *tightly coordinated* teams are those in which robots continuously interact during execution, usually to complete a task that requires multiple robots such as the joint transportation of a large object. Tight coordination is extremely challenging: teams cannot easily take advantage of the distributed planning and execution that make loose coordination tractable, and they are rarely fault-tolerant since task success depends on the simultaneous success of multiple teammates.

## 8.1 Related Work

Market-based approaches are not often used to directly facilitate tight coordination. In most domains, the interactions between tightly-coordinating robots are simple and do not warrant the extra computation and communication expenses that come with many market-based approaches. For instance, Simmons *et al.* [39] propose auctions to secure teammate participation in subtasks that require tight coordination between multiple robots, but implement the tight coordination using an inexpensive reactive approach.

However, some domains such as security sweeping require planned tight coordination and cannot be effectively solved with simpler approaches. Hoplites [23] is a market-based approach to solving such domains. The approach couples cost and revenue between teammates to facilitate coordination. Robots broadcast their intentions to teammates, locally plan interactions in easier problem scenarios, and negotiate participation in tightly-coupled joint plans in harder problem scenarios.

## 8.2 Future Challenges

Significant challenges remain in using market-based techniques for tight coordination. Among other requirements the field needs some formalization of solution quality, market mechanisms that can handle persistent coordination, and bidding techniques that concisely encapsulate complex constraints. Additionally, even in the general multi-robot literature, tight coordination in highly dynamic environments a largely remains unsolved problem.

## 9 Learning and Adaptation

While a generalized system is useful, its application to specific domains usually requires some adaptation. Online opportunistic adaptation is therefore a highly relevant and useful feature. However, in dynamic environments where teams can fluidly change in size, where interaction strategies can be continuously modified, and where external conditions can be unexpectedly changed, *a priori* definition of best trading and coordination strategies can be very difficult, and sometimes impossible. Consequently, the robots require not only the ability to quickly adapt their behaviour in response to dynamic events and to changes in the other agents' behaviour, but also the ability to determine when and how this adaptation should take place. Hence, the integration of learning techniques can be a very powerful feature.

### 9.1 Related Work

The application of learning techniques in market-based coordination is currently at a very early stage. One big debate is whether learning should be applied at the team level or at the individual level or some combination of the two. Another important question to be answered is how to deal with team interactions – should other agents be dealt with as environmental factors or should they be dealt with in a special way? Oliveira *et al.* [32] present a detailed discussion of the issues relevant to the application of learning in dynamic markets. The role that learning can play in market-based multirobot coordination is also discussed briefly by Stentz and Dias [41].

The authors are unaware of any learning techniques implemented on a team of physical robots coordinated using a market-based approach. However, publications are starting to emerge in the application of learning techniques for market-based coordination of simulated robot teams. Notably, learning techniques are applied to learn bidding strategies in dynamic markets [32], opportunity costs in a simulated distributed sensing task [37], and role assignment [26] and bidding strategies [15] in simulated robot soccer.

### 9.2 Future Challenges

The application of relevant learning techniques to market-based coordination of robot teams is a wide open research area with tremendous potential for improving team performance in dynamic environments, reducing the requirement for accuracy in cost esti-

mation and *a priori* knowledge, and enabling easy portability to different domains and environments.

## 10 Generality

Coordination approaches that are general and applicable to a variety of domains have larger overall impact. Generality requires a coordination approach to be flexible across application domains and extensible to allow easy extension of functionality and portability. Also important for generality are implementation guidelines for different domains and useful comparisons of different approaches to guide the selection of the most effective coordination scheme for a given application.

### 10.1 Related Work

#### 10.1.1 Flexibility

Since different applications will have different requirements, a widely applicable coordination approach will need to be easily configurable for the different problems it proposes to solve. Instructions and advice on how to reconfigure the mechanism for different applications will also be useful. Identifying important parameters that need to be changed based on the application requirements, instructions on how to change them, identifying components of the mechanism that need to be added/changed based on application requirements, and instructions on how to make these alterations are all important elements of a successful coordination mechanism. A further bonus will be well-designed user interfaces and tools that allow plug-and-play alterations to the coordination mechanism and automated methods for parameter tuning. The authors are aware of only two market-based approaches, MURDOCH [18] and *TraderBots* [8, 45, 44], that have been demonstrated in more than one application. However, there is much that still needs to be done in terms of providing a flexible market-based multirobot coordination approach.

#### 10.1.2 Extensibility

The ability to easily add and remove functionality is a key characteristic to building a generalized system that can evolve with the needs of the different applications. A common approach to incorporating extensibility is to build the system in a modular fashion so that different modules can be altered or replaced relatively easily according to the requirements of the specific application. In market-based approaches, it is best to modularize and isolate cost and reward functions as much as possible from task and role specifications, communication protocols, and task executives.

#### 10.1.3 Implementation

As with any claim, a proven implementation is most convincing. Moreover, successful implementation of a coordination mechanism on a robotic system requires discovering and solving many details that are not always apparent in theory, simulation and

software systems. Finally, implementation of an approach on many different platforms in a variety of application domains provides valuable insights and guidelines on how to design and implement different components of the approach in a extensible and flexible manner. Although several have been implemented on physical robot teams (e.g. [8, 18, 45, 44, 23, 39, 38]), market-based approaches have yet to be proven in a wide variety of domains.

#### 10.1.4 Comparisons

Comparisons are important to provide guidelines on how to evaluate different coordination approaches when deciding which approach is best for a given application. However, comparing different coordination approaches is a highly challenging endeavor since many considerations need to be addressed. Some of the challenges in providing a comparative framework for coordination approaches are explored by Gerkey [17] and Dias [8]. Rabideau *et al.* [33] also undertake a small comparative study that includes a market-based approach. Although market-based approaches perform well in the comparative studies they have been included in thus far, these studies are fairly limited. Thus, meaningful comparative studies of different coordination approaches remain in high demand.

### 10.2 Future Challenges

Market-based multirobot coordination approaches have only been implemented and tested in a few application domains to date. Thus, understanding and implementing generality in market-based approaches still requires significant work. However, the growing popularity of market-based methods for coordinating robot teams will be a large contributing factor to inspiring generality in this research area.

## 11 Conclusions and Future Directions

The vision that drives research in multirobot systems is that teams of robots and humans will inevitably be an integral part of our future. To realize this vision, robots must be made capable of executing complex tasks as part of a highly heterogenous team. While many multirobot coordination approaches have been proposed by the research community, market-based approaches in particular have grown in popularity over the past few years and now warrant a survey of the relevant literature. This publication meets this need by providing the first survey of the state of the art in market-based multirobot coordination approaches, thus providing three significant contributions to the multi-robot literature: an introduction to market-based multirobot coordination approaches, a review of the relevant literature, and a discussion of remaining challenges in this research area. In this survey, the authors, who significantly contributed to pioneering this research area, share their collective experience of several years in designing and implementing relevant coordination mechanisms. A detailed version of this survey is presented by Kalra *et al.* [24].

Evident in this publication is a large body of literature relevant to market-based multirobot coordination. The relevant work ranges from theoretical formulations to conceptual design frameworks to implementations in simulation and on physical robot teams. Chosen application domains span a wide range including distributed sensing, mapping, exploration, surveillance, perimeter sweeping, assembly, box-pushing, reconnaissance, and soccer. However, this is still a relatively young field of research, and many challenges still must be overcome. While challenges specific to each of the reviewed categories are presented in the relevant sections, the most significant overall challenges in this area of research are examined next.

A first important need is a clear conceptual understanding of market based coordination approaches. To that end, this paper provides a broadly inclusive definition of market-based approaches but further discussion is needed to characterise the domain and further our understanding of how components such as cost and reward functions, bidding strategies, and auction clearing mechanisms can be designed, implemented, and used effectively in different application domains. A second challenge is to demonstrate long-term, reliable, and robust operation of robot teams coordinated using market-based approaches in a variety of applications. A third category of important research is providing solution quality measures and performance guarantees.

To deal with dynamic environments, heterogeneous teams, and tasks that require tight-coordination, researchers must address important research challenges which include enabling replanning, designing appropriate penalties for contract breaches, allowing different levels of commitments to contracts, understanding how to deal with subteams more effectively, and applying appropriate learning techniques. Finally, formulating and carrying out meaningful comparisons of market-based approaches with other coordination approaches and providing guidance for selecting appropriate coordination techniques for a given application are paramount for advancing our collective understanding of multirobot coordination.

Despite the many challenges ahead, market-based techniques have shown much promise as versatile and powerful coordination schemes for groups of robots executing complex tasks. This survey and analysis is intended to review the state of the art in this area of research and provide inspiration and direction for future research efforts on this topic.

## Acknowledgments

This work is sponsored in part by the Boeing Company Grant CMU-BA-GTA-1 and in part by the U.S. Army Research Laboratory, under contract **Robotics Collaborative Technology Alliance** (contract number DAAD19-01-2-0012). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Boeing Company, the Army Research Laboratory or the U.S. Government.

## References

- [1] M. Andersson and T. Sandholm. Contract type sequencing for reallocation negotiation. In *International Conference on Distributed Computing Systems*, 2000.
- [2] K. Azarm and G. Schmidt. A decentralized approach for the conflict free motion of multiple mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1996.
- [3] C. Bererton, G. Gordon, S. Thrun, and P. Khosla. Auction mechanism design for multi-robot coordination. In *Advances in Neural Information Processing Systems*, 2003.
- [4] M. Berhault, H. Huang, P. Keskinocak, S. Koenig, W. Elmaghraby, P. Griffin, and A. Kleywegt. Robot exploration with combinatorial auctions. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2003.
- [5] S. S. C. Botelho and R. Alami. M+: A scheme for multi-robot cooperation through negotiated task allocation and achievement. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1999.
- [6] P. Caloud, W. Choi, J. C. Latombe, C. L. Pape, and M. Yim. Indoor automation with many mobile robots. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 1990.
- [7] P. Chandler and M. Pachter. Hierarchical control for autonomous teams. In *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, 2001.
- [8] M. B. Dias. *TraderBots: A New Paradigm for Robust and Efficient Multirobot Coordination in Dynamic Environments*. PhD thesis, Robotics Institute, Carnegie Mellon University, January 2004.
- [9] M. B. Dias, B. Browning, M. M. Veloso, and A. Stentz. Dynamic heterogeneous robot teams engaged in adversarial tasks. Technical Report CMU-RI-TR-05-14, Robotics Institute, Carnegie Mellon University, 2005.
- [10] M. B. Dias, D. Goldberg, and A. Stentz. Market-based multirobot coordination for complex space applications. In *the 7th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS)*, 2003.
- [11] M. B. Dias and A. Stentz. Opportunistic optimization for market-based multirobot control. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2002.
- [12] M. B. Dias, R. Zlot, M. Zinck, J. P. Gonzalez, and A. Stentz. A versatile implementation of the *TraderBots* approach to multirobot coordination. In *Proceedings of the International Conference on Intelligent Autonomous Systems (IAS)*, 2004.

- [13] M. B. Dias, R. Zlot, M. Zinck, and A. Stentz. Robust multirobot coordination in dynamic environments. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.
- [14] D. J. Farber and K. C. Larson. The structure of a distributed computing system – software. In *Proceedings of the Symposium on Computer-Communications Networks and Teletraffic*, 1972.
- [15] V. Frias-Martinez and E. Sklar. A team-based co-evolutionary approach to multi agent learning. In *Proceedings of the 2004 AAMAS Workshop on Learning and Evolution in Agent Based Systems*, 2004.
- [16] V. Frias-Martinez, E. Sklar, and S. Parsons. Exploring auction mechanisms for role assignment in teams of autonomous robots. In *Proceedings of the RoboCup Symposium*, 2004.
- [17] B. P. Gerkey. *On Multi-Robot Task Allocation*. PhD thesis, Computer Science Department, University of Southern California, August 2003. CRES-03-012.
- [18] B. P. Gerkey and M. J. Matarić. Sold!: Auction methods for multi-robot control. *IEEE Transactions on Robotics and Automation Special Issue on Multi-Robot Systems*, 18(5), 2002.
- [19] D. Goldberg, V. Ciciello, M. B. Dias, R. Simmons, S. Smith, and A. Stentz. Market-based multi-robot planning in a distributed layered architecture. In *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings from the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers, 2003.
- [20] M. Golfarelli, D. Maio, and S. Rizzi. A task-swap negotiation protocol based on the contract net paradigm. Technical Report 005-97, CSITE (Research Center for Informatics and Telecommunication Systems), University of Bologna, 1997.
- [21] J. Guerrero and G. Oliver. Multi-robot task allocation strategies using auction-like mechanisms. In *Sixth Congress of the Catalan Association for Artificial Intelligence (CCIA)*, 2003.
- [22] L. Hunsberger and B. J. Grosz. A combinatorial auction for collaborative planning. In *Proceedings of the Fourth International Conference on Multi-Agent Systems (ICMAS)*, 2000.
- [23] N. Kalra, D. Ferguson, and A. Stentz. Hoplitest: A market-based framework for complex tight coordination in multi-robot teams. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [24] N. Kalra, R. Zlot, M. B. Dias, and A. Stentz. Market-based multirobot coordination: A comprehensive survey and analysis. Technical Report CMU-RI-TR-05-16, Robotics Institute, Carnegie Mellon University, 2005.



- [25] H. Köse, Çetin Meriçli, K. Kaplan, and H. L. Akin. All bids for one and one does for all: Market-driven multi-agent collaboration in robot soccer domain. In *Proceedings of the Eighteenth International Symposium on Computer and Information Sciences*, 2003.
- [26] H. Köse, U. Tatlıdede, Çetin Meriçli, K. Kaplan, and H. L. Akin. Q-learning based market-driven multi-agent collaboration in robot soccer. In *The Turkish Symposium on Artificial Intelligence and Neural Networks*, 2004.
- [27] M. Lagoudakis, V. Markakis, D. Kempe, P. Keskinocak, S. Koenig, A. Kleywegt, C. Tovey, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Robotics: Science and Systems*, 2005.
- [28] M. G. Lagoudakis, M. Berhault, S. Koenig, P. Keskinocak, and A. J. Kleywegt. Simple auctions with performance guarantees for multi-robot task allocation. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2004.
- [29] T. Lemaire, R. Alami, and S. Lacroix. A distributed tasks allocation scheme in multi-UAV context. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2004.
- [30] D. C. MacKenzie. Collaborative tasking of tightly constrained multi-robot missions. In *Multi-Robot Systems: From Swarms to Intelligent Automata: Proceedings of the 2003 International Workshop on Multi-Robot Systems*, volume 2. Kluwer Academic Publishers, 2003.
- [31] R. Nair, T. Ito, M. Tambe, and S. Marsella. Task allocation in the rescue simulation domain: A short note. In *RoboCup-2001: The Fifth Robot World Cup Games and Conferences*. Springer-Verlag, 2002.
- [32] E. Oliveira, J. M. Fonseca, and N. R. Jennings. Learning to be competitive in the market. In *Proceedings of the AAI Workshop on Negotiation: Settling Conflicts and Identifying Opportunities*, 1999.
- [33] G. Rabideau, T. Estlin, S. Chien, and A. Barrett. A comparison of coordinated planning methods for cooperating rovers. In *Proceedings of the AIAA 1999 Space Technology Conference*, 1999.
- [34] T. Sandholm. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*, 1993.
- [35] T. Sandholm. Contract types for satisficing task allocation: I theoretical results. In *AAAI Spring Symposium: Satisficing Models*, 1998.
- [36] T. Sandholm. Algorithm for optimal winner determination in combinatorial auctions. *Artificial Intelligence*, 135(1), 2002.

- [37] J. Schneider, D. Apfelbaum, D. Bagnell, and R. Simmons. Learning opportunity costs in multi-robot market based planners. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [38] R. Simmons, D. Apfelbaum, W. Burgard, D. Fox, M. Moors, S. Thrun, and H. Younes. Coordination for multi-robot exploration and mapping. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, 2000.
- [39] R. Simmons, S. Singh, D. Hershberger, J. Ramos, and T. Smith. First results in the coordination of heterogeneous robots for large-scale assembly. In *Proceedings of the International Symposium on Experimental Robotics (ISER)*, December 2000.
- [40] R. G. Smith. The contract net protocol: High-level communication and control in a distributed problem solver. *IEEE Transactions on Computers*, 29(12), 1980.
- [41] A. Stentz and M. B. Dias. A free market architecture for coordinating multiple robots. Technical Report CMU-RI-TR-99-42, Robotics Institute, Carnegie Mellon University, December 1999.
- [42] S. Thayer, B. Digney, M. B. Dias, A. Stentz, B. Nabbe, and M. Hebert. Distributed robotic mapping of extreme environments. In *Proceedings of SPIE: Mobile Robots XV and Telemanipulator and Telepresence Technologies VII*, 2000.
- [43] C. Tovey, M. G. Lagoudakis, S. Jain, and S. Koenig. The generation of bidding rules for auction-based robot coordination. In *Proceedings of the 3rd International Multi-Robot Systems Workshop, Naval Research Laboratory*, 2005.
- [44] R. Zlot and A. Stentz. Complex task allocation for multiple robots. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [45] R. Zlot, A. Stentz, M. B. Dias, and S. Thayer. Multi-robot exploration controlled by a market economy. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2002.